

[◆ FREE GUIDE](#)

AI ESSENTIALS · GUIDE NO. 20

THE FRIDAY RECONCILIATION

Every Friday, have AI compare what you said you would finish against what you actually finished. Not from your task list, from the record: your sent mail, your docs, your calendar notes, your messages. Here is the exact system, the standing prompt, and the guardrail that keeps it safe.

~11 MIN READ · FOR EVERYONE WHO COMMENTED "FRIDAY" · FIRST SETUP: ~25 MINUTES

Hank Barker // AI Essentials // @hankhelpsai

THE SHORT VERSION

- Most of what you promised this week never became a task. It lived in a sentence: "I'll send that over," "I'll have it ready by Friday."
- The reconciliation has AI search your **own words** across email, docs, calendar notes and messages, then look for **evidence** that each promise was kept.
- Everything lands in three groups: **completed, still open, no longer relevant**.
- For anything still open it prepares the next action **and shows you the evidence it used**, so you can check its reasoning in seconds.
- Nothing goes out automatically. As of July 2026 Anthropic's Google Workspace connector **cannot send Gmail at all**, only draft, so the guardrail is built into the tool.

WHY THE TASK LIST QUIETLY LIES

A weekly review built on your task list is a review of the promises you were organized enough to write down. That is a much smaller set than the promises you actually made.

Think about the last five days. Somewhere in a thread you told a client you would send the revised numbers. In a meeting note you typed "I'll pull the Q3 figures before we meet again." On a call you agreed to loop in someone from legal, and afterwards you dropped a line about it in the calendar event so you would not forget. In a chat someone asked for the updated deck and you replied "leave it with me."

None of those are tasks. They are sentences. They got made in the flow of a conversation, at the exact moment when stopping to write them down would have interrupted the thing you were doing. That is not a discipline problem, it is just how commitments actually get made by people with real jobs.

The ones that cause damage are the ones you forgot you made. The other person did not forget. Two weeks later it comes back as "any update on that?" and now you are on the back foot for a reason that has nothing to do with the quality of your work. You did not drop the ball because you were careless. You dropped it because it was never on the list you were checking.

Reconciliation is an accounting word on purpose. You are matching two ledgers: what you said, and what the record shows actually happened.

So the fix is not a better task list. The fix is to stop reviewing your memory and start reviewing the record. You already have the record. It is sitting in your sent mail, your documents, your calendar and your messages. What you have never had is a way to read all of it at once. That is the job you can hand to AI now, and it is a genuinely good use of it: tedious, wide, evidence-based, and completely checkable.

WHERE COMMITMENTS ACTUALLY LIVE

Four surfaces hold almost everything, and it is worth naming them precisely, because the search you write is only as good as the places you point it at.

- **Email.** Mostly your sent folder, plus your own replies inside threads other people started. This is where the majority of commitments to people outside your team get made.
- **Documents.** Meeting notes, action item lists, project docs. Look at what you edited this week, not what exists.
- **Calendar notes.** The description field, and whatever you typed into an event after a call ended. Underrated, and it is where the vaguest commitments hide.
- **Messages.** Team chat and DMs, where commitments are shortest and least likely to be written down anywhere else.

Now the part that makes this work. You do not search for topics, you search for **language**. A commitment has a grammatical shape: first person, future tense, an object, and often a deadline. "I'll send the deck by Thursday" is one. "The deck is interesting" is not. When you tell the AI to look for the phrasing rather than the subject matter, it stops needing to understand your projects and starts recognizing a pattern it can find reliably.

COMMITMENT LANGUAGE: SEARCH FOR PHRASES, NOT TOPICS

"I'll send"	"I'll update"	"I'll get you"	"I'll circle back"
"I'll follow up"	"let me pull"	"I can have that by"	"ready by Friday"
"by end of week"	"leave it with me"	"I'll take that one"	"I'll look into"
"I'll draft"	"I'll confirm"	"I'll put together"	"I'll loop in"

These are your own words, in your sent mail, your notes, and your messages.
Add the three or four phrases you personally overuse. That is the whole tuning job.

Search for the shape of a promise, not the topic of one. This list is the seed; the useful version has the phrases you personally overuse added to it.

Two boundaries keep the list clean. Soft language is not a commitment: "I'll think about it," "we could probably," "happy to take a look sometime." Those go nowhere. And commitments other people made *to you* are a different list with a different purpose, so keep them out of this one or the review turns into a grievance log.

SCOPE IT TO YOUR OWN WORDS

The review only covers things **you** said you would do. It is a self-audit, not a chase list. If you want the other direction later, that is a second scheduled task with the search reversed, and it reads very differently.

THE STEP THAT MAKES IT HONEST

If you stop after finding the commitments, you have produced a list of things you may or may not have done, and you are right back to checking it against your memory. The step that turns this from a nice idea into a real system is evidence. For every single promise, the AI has to go looking for proof that it was completed.

FOUR THINGS THAT COUNT AS PROOF**A sent reply**

The thing was delivered, and the timestamp proves when.

**An updated file**

A doc or sheet changed after the promise was made.

**A closed task**

Something was ticked off in a tracker or project tool.

**A later message**

The agreement moved, shrank, or was called off.

Every item in the finished review has to point at one of these, quoted.

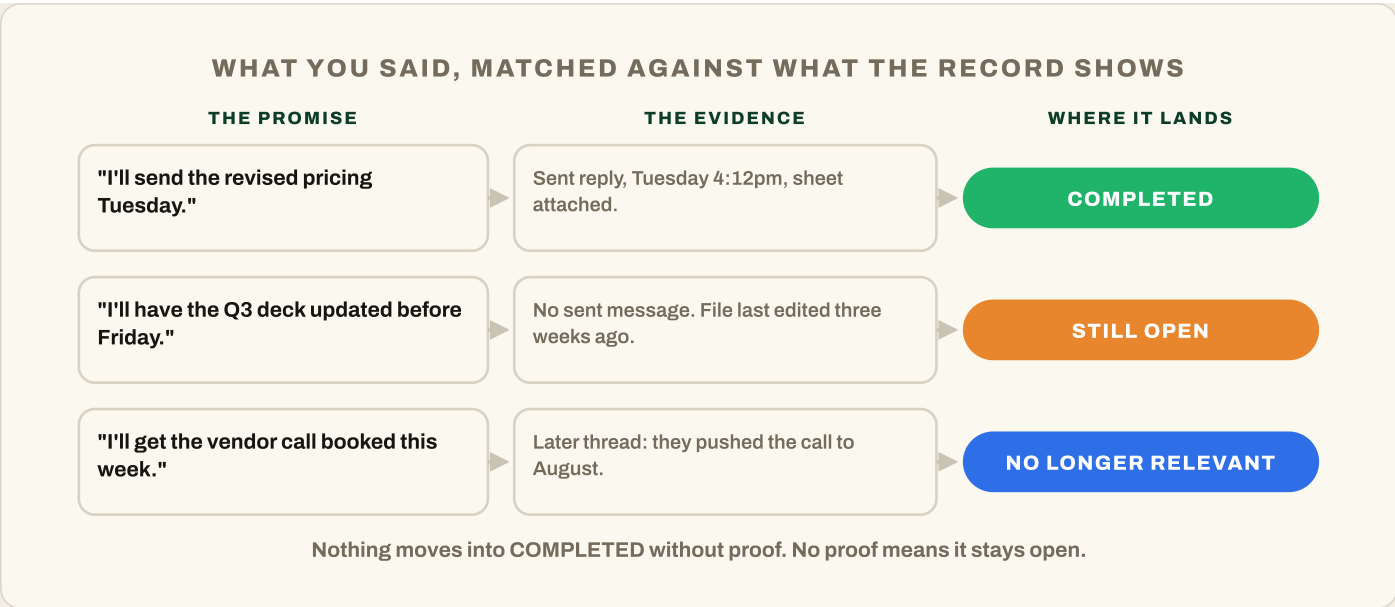
Four kinds of proof, and the last one matters as much as the first three. Agreements change, and a changed agreement is not a failure.

The fourth kind is the one people forget to ask for, and it is the reason this review does not feel like a weekly guilt trip. A large share of what looks unfinished was actually renegotiated. The client pushed the date. The project got cut. Someone else picked it up. A later message in the same thread quietly replaced the original agreement, and if the AI reads only the promise and not the rest of the conversation, it will report a failure that never happened.

One rule keeps the whole thing trustworthy: absence of evidence is not evidence of completion. If the AI cannot find proof, the item stays open and it says plainly that it found nothing. That bias costs you a few false alarms and it saves you from the only failure mode that actually hurts, which is being told something shipped when it did not.

THREE GROUPS, NOTHING ELSE

The finished review sorts into exactly three buckets. Resist adding more. The value here is that a week of scattered promises collapses into one short page you can read standing up.



The whole system in one picture: a promise column, an evidence column, and three places an item can land.

GROUP	WHAT IT MEANS	WHAT YOU GET
Completed	Evidence was found that it happened	One line each: the promise, the proof, the date. Read it and move on.
Still open	No evidence found, and the agreement was never changed	The promise, who is waiting, how long it has been open, the next action, and a draft ready for you to review.
No longer relevant	A later message changed or cancelled the agreement	What changed, quoted, so you can confirm you are not letting yourself off the hook.

Still open is the only group that needs work from you, so it should sit at the top of the output and be sorted oldest first. Completed exists to be scanned in fifteen seconds and to give you an honest picture of a week that probably went better than you think. No longer relevant exists so that the open list stays short and believable.

ASK FOR A FOURTH BUCKET YOU HOPE STAYS EMPTY

Tell it that anything it is genuinely unsure about goes in a **Check this** group with a one-line reason. Ambiguity parked in its own bucket is useful. Ambiguity guessed at and filed under Completed is not.

WHAT YOU NEED BEFORE YOU BUILD

The version of this that runs without you is a scheduled task in **Claude Cowork**. Cowork is Anthropic's agentic mode for knowledge work: you describe an outcome, walk away, and come back to finished work. As of July 2026 it is available in the Claude desktop app on macOS and Windows, on the web at claude.ai, and on mobile, and it requires a paid plan (Pro, Max, Team or Enterprise). You start it by opening Claude and choosing **Cowork** rather than Chat in the message box.

The screenshot shows a web browser window with the URL `support.claude.com/en/articles/13345190-get-started-with-claude-cowork`. The page header includes the Claude Support logo, navigation links for API Docs, Release Notes, How to Get Support, and a language selector set to English. A search bar is present below the header.

The article title is "Get started with Claude Cowork", updated over 2 weeks ago. The main text explains how to use Claude Cowork, which brings Claude Code's agentic capabilities to knowledge work beyond coding.

Availability

Claude Cowork is available for paid plans (Pro, Max, Team, Enterprise) on:

- **Claude Desktop for macOS** — [Click here](#) to download
- **Claude Desktop for Windows** — Cowork requires the latest version of Claude for Windows. Download or update at claude.com/download.
- **Web**, at claude.ai
- **Claude Mobile**, in the latest version of Claude for iOS and Claude for Android

Chat and Cowork share one home, so you start both from the same place. On any surface, find the message box and select "Cowork," then describe your task. To go back to a regular conversation, select "Chat."

Claude Cowork is in beta on web and mobile, and rolling out over the next several weeks starting with the Max plan, with more plans to follow.

Table of Contents:

- Availability
- What is Claude Cowork?
- How Claude Cowork runs your tasks
- Get started
- Start a Cowork session
- What to expect during a task
- Choose how Claude checks with you
- Add global and folder instructions
- Claude Cowork plugins
- Schedule recurring tasks
- Usage limits
- Permissions and security
- Example use cases
- Current limitations
- Troubleshooting

Cowork is the mode where you hand over an outcome instead of a question. Everything in this guide runs inside it. — **Source:** [Anthropic, Get started with Claude Cowork](#)

Then you need connectors, because a reconciliation with nothing to read is just a prompt. The Google Workspace connector is the one that does most of the work here. As of July 2026 it lets Claude read, search and draft in Gmail and manage labels and threads; create, update and delete Google Calendar events and find mutual availability; and search and retrieve Google Docs while reading Sheets, Slides, PDFs and Office files from Drive. You turn it on from the plus sign in chat, hover **Connectors**, and toggle it. Team and Enterprise admins enable it at **Organization settings, Connectors**, and everything is managed at claude.ai/customize/connectors.


[API Docs](#)
[Release Notes](#)
[How to Get Support](#)
[English](#)

 Search for articles...

[All Collections](#) > [Connectors](#) > [Pre-Built Connectors](#) > [Use Google Workspace connectors](#)

Use Google Workspace connectors

May 22, 2026

Connect your Gmail, Google Calendar, and Google Drive to Claude so you can search emails, manage your calendar, work with documents, and save files—all without leaving the conversation.

Google Workspace connectors (Gmail, Google Calendar, and Google Drive) are available for all users on Claude and Claude Desktop.

For Team and Enterprise plans, an Owner or Primary Owner must enable these connectors at the organization level before individual users can authenticate. For setup instructions, read [Use connectors to extend Claude's capabilities](#).

What you can do

Gmail

- **Search and read emails** using natural language queries
- **Draft emails** with proper formatting and context — Claude creates drafts in your Gmail account, but cannot send emails on your behalf
- **Access email metadata**, including attachment metadata (not attachment content)

What you can do

- How to use the connectors
- Privacy and data handling
- Current limitations
- Troubleshooting
- Frequently asked questions

Gmail, Calendar and Drive in one connector. Note what it says about Gmail, because that limitation turns out to be the best feature in this whole build. — **Source:** [Anthropic, Use Google Workspace connectors](#)

If your notes live in Notion, connect it from Claude's connector directory rather than wiring up a raw MCP URL. Worth knowing before you click approve: Notion's official connection authorizes over OAuth and grants access to your workspace at your own permission level, so it can read and write to your pages just like you can. That is fine for this job, and it is still worth knowing.

YOU MAY NOT SEE THE SCHEDULED TAB YET

Anthropic's own page says Claude Cowork is in beta on web and mobile, and is rolling out over several weeks **starting with the Max plan**, with more plans to follow. So if you are on Pro and the Scheduled section is not there, you have not done anything wrong. Build the prompt now, run it manually this Friday, and schedule it when the tab appears.

BUILD THE FRIDAY TASK

Scheduled tasks live under **Scheduled** in the Cowork left sidebar. From there you choose **New task**, then either **Create with Claude**, which walks you through a few multiple-choice questions before you click Schedule, or **Set up manually**, where you fill in the task name, the prompt, the approval mode and the frequency, with optional model and folder settings, and then hit Save.

The screenshot shows the Claude Support website. The URL in the browser is `support.claude.com/en/articles/13854387-schedule-recurring-tasks-in-claude-cowork`. The page title is "Schedule recurring tasks in Claude Cowork". It includes a search bar, navigation links for API Docs, Release Notes, How to Get Support, and a language dropdown set to English. The article content states that scheduled tasks allow delegating work to Claude Cowork on a recurring basis. A callout box confirms that scheduled tasks are available for all paid plans (Pro, Max, Team, Enterprise) and that the beta rollout starts with the Max plan. A sidebar on the right lists topics: What scheduled tasks can do, How scheduled tasks work, Create a scheduled task, and Manage your scheduled tasks. The main text under "What scheduled tasks can do" mentions access to capabilities like connected tools, skills, and installed plugins.

Anthropic's page for scheduled tasks, where a good prompt stops being something you remember to run. Note the box confirming scheduled tasks are available on all paid plans, and the beta rollout starting with Max. — **Source:** [Anthropic, Schedule recurring tasks in Claude Cowork](#)

1. Open Cowork and go to **Scheduled** in the left sidebar, then **New task**.
2. Choose **Set up manually**, since you already know exactly what you want it to do.
3. Name it something you will recognize in a list a month from now. "Friday reconciliation" is fine.
4. Paste the standing prompt from the next section into the prompt field.
5. Set the frequency to **weekly**. As of July 2026 the available cadences are hourly, daily, weekly, on weekdays, or manually, and a Friday review is the weekly one.
6. Set the approval mode deliberately. This is your leash, and the next section explains where to hold it.
7. Save it, then run it once by hand before you trust the schedule.

 support.claude.com/en/articles/13854387-schedule-recurring-tasks-in-claude-cowork

What scheduled tasks can do

Scheduled tasks have access to the same capabilities as regular Cowork tasks, including connected tools, skills, and installed plugins. Common uses include:

- **Daily briefings:** Summarize Slack messages, emails, or calendar events from the past 24 hours.
- **Weekly reports:** Compile data from Google Drive, spreadsheets, or connected tools into a formatted summary.
- **Recurring research:** Track topics, competitors, or industry news on a regular cadence.
- **File organization:** Periodically sort, clean up, or process files in a designated folder.
- **Team updates:** Generate status updates or standup summaries from project management tools.

What scheduled tasks can do

How scheduled tasks work

Create a scheduled task

Manage your scheduled tasks

How scheduled tasks work

When you create a scheduled task, Claude saves your prompt as the task's instructions and runs them at the cadence you choose. Tasks can search Slack, query files, run web research, generate reports, and more—using any connectors and plugins you've set up in Cowork.

Each scheduled task runs as its own Cowork session. You can review the results when they're ready, just like any other task.

Scheduled tasks run remotely, so they run on their cadence even when your computer is asleep or the Claude Desktop app is closed. Review upcoming and past runs by clicking "Scheduled" in the left sidebar on any surface.

Note: Scheduled tasks use the built-in schedule options and work with your connectors and the files saved to your Claude account. They can't be tied to a folder on your computer.

For Team and Enterprise organizations, admins control Cowork access through the admin toggle. For more details, see [Use Claude Cowork on Team and Enterprise plans](#).

How a scheduled task actually runs: Claude saves your prompt as the task's instructions, each run happens as its own fresh Cowork session, and it runs remotely on its cadence. Note the box about connectors and folders. — **Source:** [Anthropic, Schedule recurring tasks in Claude Cowork](#)

On timing, I would put it in the middle of Friday afternoon rather than at the end of the day. Around 3:00 or 4:00 PM the week is essentially over, but you are still at your desk and still willing to fire off two replies. A review that lands at 6:00 PM on Friday gets read on Monday, and by Monday it is describing a week you have already stopped thinking about.

IT RUNS WHILE YOU SLEEP, WITH ONE EXCEPTION

Anthropic's wording is that scheduled tasks run remotely, so they run on their cadence even when your computer is asleep or the Claude Desktop app is closed. The exception matters: **if a scheduled task requires local files or apps, it will only run locally.** So keep this review pointed at connected cloud services. The moment you make it read a folder on your laptop, it stops being a task that runs on its own.

THE STANDING PROMPT

This is the whole system written out. Paste both halves together as one prompt, then edit the bracketed parts to match the tools you actually use. It is specific on purpose: which surfaces to search, what counts as evidence, and exactly how the output should be shaped.

You are running my Friday reconciliation. Today is Friday.
Cover the last 7 days.

STEP 1: FIND WHAT I PROMISED

Search these surfaces for things I said I would do:

- Gmail: messages I sent, and my own replies inside threads
- Google Calendar: event descriptions and notes I added after calls
- Google Drive: docs I edited this week, especially meeting notes
- [your chat tool, Notion, or project tool]

Look for commitment language in my own words, not task titles:

"I'll send", "I'll update", "I'll get you", "I'll circle back",
"let me pull", "I can have that by", "ready by Friday",
"by end of week", "leave it with me", "I'll follow up".

For each one, capture: what I promised, who I promised it to,
where I said it, the date, and any deadline I named.

Ignore soft language like "I'll think about it".

Ignore commitments other people made to me.

STEP 2: LOOK FOR EVIDENCE

For each commitment, search for proof it was finished. Count any of:

- a sent message or reply that delivers the thing
- a file created or edited after the promise that matches it
- a task or issue marked done
- a later message where we changed or cancelled the agreement

If you find nothing, say so plainly. Never assume it was done.

STEP 3: SORT INTO THREE GROUPS

COMPLETED: evidence found. One line each: promise, proof, date.

STILL OPEN: no evidence. For each, give me the promise, who is waiting, how many days it has been open, the exact next action, and a draft ready for me to review. Save email drafts in Gmail.

NO LONGER RELEVANT: superseded or cancelled. Quote what changed.

FORMAT

Open with one line: X completed, Y still open, Z no longer relevant.

Then the three groups, still open first, oldest first inside it.

Under every item, show the evidence you used as a short quote plus a link to where you found it.

If you are unsure whether something is a commitment, put it in a fourth group called CHECK THIS and tell me why you were unsure.

Do not send any email or message. Do not create, move or delete calendar events. Drafts only. I will send everything myself.

WHY THE LAST PARAGRAPH IS NOT OPTIONAL

Scheduled task runs are fresh sessions with access to your connected tools, so the instruction has to travel inside the prompt itself. The calendar line matters more than people expect, because the Google Workspace connector has full write access to Calendar as of July 2026: it can create, update and delete events. Gmail cannot send. Calendar can act. Write the rule down.

WHY IT MUST NOT SEND ANYTHING

It is tempting to close the loop and let the review fire off its own follow-ups. Every open item already has a draft attached, so why not just let it go? Because the AI can be confidently wrong about what you actually committed to, and the failure is embarrassing in a specific way. It reads "I'll send the pricing once legal signs off" and drafts a note that sends the pricing. It reads a promise you made to a colleague as a promise you made to a client. The words were there; the meaning was not.

So the human step is not friction you should try to remove. Confirming that it understood the commitment correctly is the entire point of the review, and it takes about eight seconds per item because the evidence is sitting right underneath.

THE TOOL ENFORCES THIS FOR YOU


As of July 2026, Anthropic's Google Workspace connector states plainly that **Claude creates drafts in your Gmail account, but cannot send emails on your behalf**. Not a setting, not a toggle you might leave off by accident. The connector has no send capability at all. The most important guardrail in this entire system is structurally impossible to disable, which is exactly what you want on a task that runs unattended every week.

Two honest caveats. First, that limit is specific to this connector, so if your organization runs Microsoft 365 rather than Google Workspace, check what your own connector is permitted to do before you lean on draft-only behavior as your safety net. Second, the same protection does not extend to your calendar, which the connector can write to freely, so the "do not create, move or delete events" line in the prompt is doing real work.

WORKING THE REVIEW ON FRIDAY AFTERNOON

The output is designed to be worked, not admired. Ten minutes is the target, and the order matters.

1. Read the one-line summary. Three numbers tell you what kind of week it was before you read anything else.
2. Go straight to **Still open**, oldest first. Age is the best proxy for risk you have.
3. For each item, read the quoted evidence before the draft. You are checking whether it understood the commitment, not whether it wrote nicely.
4. Approve, edit, or reject. Rejecting is a real outcome, and about one item in ten deserves it.
5. Send the drafts yourself, from your own inbox, in your own words.
6. Skim **No longer relevant** and make sure you agree that each one really was renegotiated rather than quietly abandoned.
7. Scan **Completed** last, quickly. It is there to close the week honestly, and most weeks it is longer than you expected.



Promises accumulate all week without being written down anywhere. Friday afternoon the task runs on its own, and your only job is ten minutes of judgment.

What you are buying is not extra output. It is the end of Sunday-evening dread, the specific feeling of knowing you have forgotten something without being able to name it. Once your open list comes from the record instead of your memory, that feeling has nowhere to live.

TUNING IT OVER THE FIRST MONTH

The first run will be mediocre. That is normal, and it is fixable in four short passes. Each week, change one thing.

- **Week one, the phrases.** Read what it caught and what it missed, then add the three or four phrasings you personally use that are not on the list. Everyone has a tell. Mine is "let me get that over to you."
- **Week two, the surfaces.** If a whole category of promise is missing, it is usually because you never pointed it at the place those promises get made. Add the surface, do not add more phrases.
- **Week three, the evidence rules.** If it is marking things complete that are not, tighten what counts as proof. A reply in the thread is not the same as the attachment actually going out.
- **Week four, the noise.** Whatever keeps landing in **Check this** tells you exactly which rule is ambiguous. Fix that one rule and the bucket empties.

After a month it should be stable, and you can mostly stop thinking about it. If you want one more improvement after that, ask it to compare this week's still-open list against last week's and flag anything appearing for the third time in a row. Items that survive three reviews are almost never forgotten. They are being avoided, and that is a different problem worth seeing clearly.

IF YOU'RE TECHNICAL: WHY ROUTINES ARE WRONG HERE

Developers reading this will reach for Claude Code routines instead, and it is worth explaining why that is the wrong instinct for this particular workflow even though it is the right tool for plenty of others.

code.claude.com/docs/en/routines

Claude Code Docs English

Search... Ctrl K Ask Assistant Claude Developer Platform Claude Code on the Web

Getting started Build with Claude Code Administration Configuration Reference Agent SDK What's New Resources

Common workflows

Prompt library

Best practices

Platforms and integrations

Overview

Remote Control

Claude Code on the web

Get started

Reference

Routines

Plan in the cloud

Ultrareview

Claude Code on desktop

Mobile

Chrome extension

Computer use (preview)

Visual Studio Code

JetBrains IDEs

Code review & CI/CD

Claude Code in Slack

CLAUDE CODE ON THE WEB

Automate work with routines

Put Claude Code on autopilot. Define routines that run on a schedule, trigger on API calls, or react to GitHub events from Anthropic-managed cloud infrastructure.

Routines are in research preview. Behavior, limits, and the API surface may change.

A routine is a saved Claude Code configuration: a prompt, one or more repositories, and a set of **connectors**, packaged once and run automatically. Routines execute on Anthropic-managed cloud infrastructure, so they keep working when your laptop is closed.

Each routine can have one or more triggers attached to it:

- Scheduled:** run on a recurring cadence like hourly, nightly, or weekly, or once at a specific future time
- API:** trigger on demand by sending an HTTP POST to a per-routine endpoint with a bearer token
- GitHub:** run automatically in response to repository events such as pull requests or releases

A single routine can combine triggers. For example, a PR review routine can run nightly, trigger from a deploy script, and also react to every new PR.

Routines are available. Ask a question... Ctrl+I

Routines are available. Create and manage them at claude.ai/code/routines, or from the CLI with **de Code on the web** enabled.

On this page

Example use cases

Create a routine

- Create from the web
- Create from the CLI

Configure triggers

- Add a schedule trigger
 - Schedule a one-off run
- Add an API trigger
 - Trigger a routine
 - API reference
- Add a GitHub trigger
 - Supported events
 - Filter pull requests
 - How sessions map to events

Manage routines

- View and interact with runs
- Edit and control routines
- Repositories and branch permissions
- Connectors
- Environments and network access
- Usage and limits
- Troubleshooting

Routines run in Anthropic's cloud on a schedule, which sounds like exactly what this needs. Read the safety note on the same page before you commit. — **Source:** [Anthropic, Claude Code routines](#)

A routine is a saved prompt plus one or more **GitHub repositories** plus connectors, running on Anthropic-managed cloud infrastructure. As of July 2026 it is in research preview, it triggers on a schedule with a minimum interval of one hour, or via an API call, or on GitHub pull request and release events. That repository requirement alone rules it out for most business professionals, since a weekly commitment review has no code in it and no reason to live in a repo.

The deeper problem is the execution model. Anthropic's own documentation states that **Claude can use every tool from an included connector, including writes, without asking for permission during a run.**

Routines run autonomously with no approval prompts, and every connected connector is included by default unless you remove it. That is a reasonable design for automated code work. It is a poor fit for a workflow whose entire value is a human confirming that the AI understood a commitment before anything leaves your name.

	COWORK SCHEDULED TASK	CLAUDE CODE CLOUD ROUTINE
Requires a GitHub repo	No	Yes
Cadence floor	Hourly, daily, weekly, weekdays	Hourly minimum
Approval during a run	Approval mode is part of the task setup	None. It runs autonomously
Connectors included	The ones you connect	All of them by default, writes included
Right for this job	Yes	No, for the reason directly above

There is also a per-account daily cap on routine runs, visible in your usage settings, with one-off runs exempt. None of this makes routines bad. It makes them a build tool rather than a review tool, and a review is what you are building.

START THIS FRIDAY

THE 25-MINUTE VERSION

Connect Google Workspace. Open Cowork, paste both halves of the standing prompt, and run it once by hand right now against the last seven days. Read what comes back, add the two phrases it missed, then save it as a weekly scheduled task for Friday at 3:00 PM. Next week it shows up without you.

THE ONE THING NOT TO CHANGE

Whatever else you tune, keep the no-send rule and keep the evidence quoted under every item. Those two lines are what separate a review you can trust from a very confident weekly list of things that may or may not be true.

The whole idea is small enough to say in one sentence. Stop reviewing your memory. Review the record, on a schedule, with the proof attached, and hold on to the ten minutes of judgment that only you can supply.

SOURCES & WHERE TO LEARN MORE

- [Anthropic, Get started with Claude Cowork](#)
- [Anthropic, Schedule recurring tasks in Claude Cowork](#)
- [Anthropic, Use Google Workspace connectors \(Gmail is draft only\)](#)
- [Anthropic, Claude Code routines \(the developer path\)](#)
- [Manage your connectors](#)
- [Notion, official MCP and connector docs](#)
- [Want the connectors playbook? Comment "SYSTEM" for that guide](#)

Built by Hank Barker // AI Essentials // @hankhelpsai. Every capability claim here was checked against Anthropic's official documentation in July 2026. Cowork is in beta and moving fast, so re-check the scheduling and connector pages before you rely on a detail.